

A SWAR Approach to Counting Ones (Preliminary Report)

Holger Petersen
Reinsbugstr. 75
70197 Stuttgart
Germany

August 22, 2011

Abstract

We investigate the complexity of bit counting algorithms in different sets of instructions.

1 Introduction

The operation of counting the number of ones in a binary word consisting of n bits is also known as sideways addition [7] or bit count [6]. Solutions based on look-up tables have been suggested by several authors (see the references in [8]), but for large n they are clearly impractical.

Under unit cost measure, methods of asymptotical complexity $O(\log \log n)$ are known, but they are based on multiplication like the one due to Wilkes, Wheeler, and Gill from the first textbook on programming published in 1957 [7] or division as in [3, Item 169]. The assumption that these complex arithmetical operations are executed as efficiently as, e.g., addition or logical instructions appears to be unrealistic. We are therefore mainly interested in approaches to the bit count problem avoiding these operations, also ruling out shift instructions as special cases of multiplication and division.

Exercise 2-9 of [6] asks for a solution of the bit count problem based on an observation originally due to Wegner [9]. It is a remarkable fact that in two's complement representation for $x \neq 0$ the right-most one of x can be deleted by the operation $x \& (x - 1)$, where $\&$ denotes the bit-wise and of two values. With the help of this fact the bit count of the input can be computed in a loop that is executed once for each one. This will save considerable work in comparison to a naive implementation when the number of ones is sparse. When the zeroes are sparse the method can be applied to the complement (this was already mentioned in the last section of [9]). The worst case as well as the average case complexity of these methods is however $\Theta(n)$.

In the current work we will develop algorithms for counting the ones in a binary word based on the concept of “SIMD with a register (SWAR)” [5]. The idea is to partition large registers into fields of bits that are processed in parallel. We obtain algorithms of complexity $O(\sqrt{n})$ and $O(\log^2 n)$ in slightly different

settings, where the latter approach is less practical for current word-lengths, as it requires a lot of large constants.

Computing the parity of a binary word efficiently has received some attention in its own right, see, e.g., [1, 4]. It can of course be determined from the count of ones, but a specific approach might be faster. For a modified parity function we get an $O(\log n)$ solution in a restricted computational model which is competitive with methods using shift instructions. Also included is a parity function making use of integer division which for 32 bits is superior (in terms of “C-operations”) to the implementations the author is aware of. It is inspired by [3, Item 169], which we outline in the Appendix.

2 Preliminaries

We will denote the binary logarithm by \log . The word-length will be denoted by n , which is usually a power of 2. Bit positions are numbered starting at 0 with the least significant bit. Thus bits of the words being processed have weights 1 to 2^{n-1} . Using the notation from [7], the bit count function is called ν .

The main results refer to programming models that rule out multiplication, division, and shift instructions. The restricted set of instructions including only logical operations (and, or, xor) and addition which will be called here Oblivious Parallel Addition and Logical Instructions (OPAL). Including subtraction does not change this model in terms of asymptotic complexity, since it can be simulated with negation and addition. OPAL is the model employed in [8]. Code will be presented in (subsets) of C [6].

With Parallel Addition and Logical Instructions (PAL) we will denote the extension by flow control instructions like “if ” and “while”, usually found in modern programming languages. Within the PAL model the bit count function from [6] and the improvement of Exercise 2-9 can be realized directly.

3 Results

Theorem 1 *The bit count function ν can be computed with the help of $O(\sqrt{n})$ instructions in the PAL model.*

Proof: The central idea is to apply Wegner’s technique [9] to approximately \sqrt{n} fields of \sqrt{n} bits each, separated by spacer bits [5]. Since the original input does not have spacer bits, the first task is to count ones at the positions of the prospective spacer bits. This number is stored in the variable `sum`.

Then the following steps are repeated. First the spacer bits are set. By subtracting twice the current most significant bits of each field (plus one for the least significant position), a one is deleted from all fields that contain a one. For fields without a one, the spacer bit is reset by this operation. The difference to the previous state of the most significant bits of the fields is computed and the change is determined with the help of the same technique. The current number of “active” fields is stored in the variable `count` and additional bits reset are recorded in `sum`. The subtraction from an empty field will cause a borrow and therefore no explicit subtraction from the next field is necessary. These steps are carried out until all fields are empty.

The number of iterations is bounded above by the length of the fields and thus $O(\sqrt{n})$. The amortized cost of updating variable `count` is also $O(\sqrt{n})$, since each of the $O(\sqrt{n})$ spacer bits can change only once (it will never change from 0 to 1 within the loop). Therefore the claimed complexity follows.

We present the procedure just outlined for an 32 bit input, where for the sake of simplicity we round the length and the number of fields to powers of 2:

```
#define HIBITSL 0x88888888 // 8 fields with 4 bits each
bitcount(a) // ALU-based
register unsigned long a;
{
    register int sum;           // overall bit sum
    register int count;        // incremental contributions
    register unsigned long x; // auxiliary variable
    register unsigned long oldhi, newhi; // hi bits of fields

    x = (a & HIBITSL);
    for(sum = 0; x ; x &= (x-1))
        sum++;
    count = 8; // bitcount(HIBITSL)
    oldhi = HIBITSL;
    a |= oldhi;
    while(oldhi)
    {
        a &= (a - oldhi - oldhi - 1);
        newhi = (a & oldhi);
        x = newhi ^ oldhi;
        while(x) {
            x &= (x-1);
            count--;
        }
        oldhi = newhi;
        sum += count;
    };
    return(sum);
}
```

□

Next we exclude flow control (`if`, `while`) from the operations allowed. By the result from [8], function ν cannot be computed under this restriction. We can however compute a modified version of ν for a portion of the bits.

First we describe an important building block that allows us to shift bits rapidly using only addition and logical operations.

Lemma 1 *In a set of disjoint fields, the least significant bit of each field can be shifted to the most significant position in constant time, if all intermediate positions contain zero.*

Proof: We will describe the technique for a single field stretching from position i to position $j > i$. By incorporating information for the other fields into the constants, the corresponding modifications can be achieved.

First bit j is cleared by masking with $\neg 2^j$. Then the value $2^j - 2$ is added to the result. Finally we mask with $\neg(2^j - 2)$.

If bit i initially was 1, then there will be a carry that propagates to position j . If it was 0, then position j will not be modified. \square

We note that the technique of Lemma 1 can replace the bit-by-bit shift of bits in the construction from [8], which has a time complexity proportional to the number of positions.

Theorem 2 *Let $m = n - \lceil \log n \rceil$. The modified bit count function $2^{m-1} \nu(x \bmod 2^m)$ for the lowest m bits of an n bit input x can be computed with the help of $O(\log^2(n))$ instructions in the OPAL model.*

Proof: We will demonstrate how to compute the function for the lower $n/2$ bits of a word with n bits. The claim then follows by computing in addition the function for the most significant $n/2 - \lceil \log n \rceil$ bits (shift all constants accordingly), adding the results and subtracting the count of ones in the overlap by first shifting each bit to the target area and then subtracting it. The latter correction can be done in time $O(\log n)$ by Lemma 1.

In stage i of the method, counts of ones in fields of length 2^{i-1} are combined into counts for fields of length 2^i by applying Lemma 1 to the bits of the counts. Since by Lemma 1 we can only shift left, the counts are located at the most significant bit of fields.

There are $O(\log n)$ stages and each stage takes time $O(\log n)$, resulting in the claimed bound $O(\log^2(n))$.

The following code implements the method for the lower 16 bits of a 32 bit word. Binary representations of the masks are given in the comments

```
// least significant 16 bits counted, 20 bits required
// (excess 4 bits for 5 bit count)
bitcount(x)
register unsigned x;
{
    register unsigned y;

    y = x & 0x5555; // 0101010101010101
    x -= y;         // delete fields
    y += 0x5555;    // 0101010101010101
    y &= 0xAAAA;    // 1010101010101010
    x += y;

    y = x & 0x6666; // 0000110011001100110
    x -= y;         // delete fields
    y += 0xCCCC;    // 0001100110011001100
    y &= 0x7333;    // 1110011001100110011
    y += 0x6666;    // 0000110011001100110
    y &= 0x7999;    // 1111001100110011001
    x += y;

    y = x & 0x3838; // 0000011100000111000
    x -= y;         // delete fields
```

```

    y += 0x1E1E0;      // 0011110000111100000
    y &= 0x61E1F;      // 1100001111000011111
    y += 0xF0F0;       // 0001111000011110000
    y &= 0x70F0F;      // 1110000111100001111
    y += 0x7878;       // 0000111100001111000
    y &= 0x78787;      // 1111000011110000111
    x += y;

    y = x & 0x780;     // 0000000000000011110000000
    x -= y;             // delete fields
    y += 0x3FC00;      // 0000000111111110000000000
    y &= 0x7C03FF;     // 0011111000000001111111111
    y += 0x1FE00;      // 0000000011111111000000000
    y &= 0x3E01FF ;    // 0001111100000000111111111
    y += 0xFF00;       // 0000000001111111100000000
    y &= 0x1FF00FF;    // 1111111110000000011111111
    y += 0x7F80;       // 0000000000111111110000000
    y &= 0xF807F;      // 0000011111000000001111111
    x += y;

    return(x);
}

```

□

By some pre- and post-processing we can obtain a solution in the PAL model of the same asymptotic complexity. Due to the large constant factor it does however not appear to be of practical value for small word-lengths.

Corollary 1 *The modified bit count function $\nu(x)$ of an n bit input x can be computed with the help of $O(\log^2(n))$ instructions in the PAL model.*

Proof: We proceed as follows:

1. The $\Theta(\log n)$ most significant positions not handled by the method of Theorem 2 are copied to a variable and set to 0 in x (time complexity $O(1)$).
2. For the $n - \Theta(\log n)$ bits the modified bit count function is computed according to Theorem 2 (time complexity $O(\log^2 n)$).
3. The resulting $O(\log n)$ bits are transferred one by one to the least significant positions by extracting each bit with the help of a mask and building up the result in another variable (time complexity $O(\log n)$).
4. The most significant bits of the original input saved in step 1) are counted in a naive way and added to the result of step 3) (time complexity $O(\log n)$).

The overall time complexity is dominated by step 2) and thus $O(\log^2 n)$. □

For the parity function only a single bit has to be shifted in an approach as in Theorem 2. Therefore we obtain a more efficient solution than for bit count.

Theorem 3 *The modified parity function $2^{n-1}(\nu(x) \bmod 2)$ can be computed with the help of $O(\log n)$ instructions in the OPAL model.*

Proof: In the initial stage the bits of the input are moved up one position and the parity of pairs of bits is computed by a XOR. By masking out the lower bits of each pair, we obtain 2 bit fields containing the parity in the most significant bit.

In the following stages we apply a variant of Lemma 1, where we propagate the most significant bit of one field directly to the most significant bit of the next field. We make use of the fact that in the least significant positions of each field spacer bits in the sense of [5] are available after masking the most significant bit. Each stage doubles the size of the fields.

We illustrate the method for 32 bits in the following code:

```
parity(x)
register unsigned x;
{
    x ^= (x + x);
    x &= 0xaaaaaaaa; // parity of 2 bit fields
    x += 0x66666666;
    x &= 0x88888888; // 4 bit fields
    x += 0x78787878;
    x &= 0x80808080; // 8 bit fields
    x += 0x7f807f80;
    x &= 0x80008000; // 16 bit fields
    x += 0x7fff8000;
    x &= 0x80000000; // all 32 bits
    return(x);
}
```

□

By the result from [8] the OPAL model is not able to directly compute the parity function. With the help of a single test a non-zero result can however be transformed into 1. We thus obtain:

Corollary 2 *The parity function $\nu(x) \bmod 2$ can be computed with the help of $O(\log n)$ instructions in the PAL model.*

We finally include code for computing the parity function inspired by the bit count function from [3] of asymptotic complexity $O(\log \log n)$. The novel approach is to xor bits first and then count the resulting ones. By reducing the number of possible ones, the fields can be smaller than in the more general setting of a full bit count.

More precisely, if we choose a field size of k bits, then it suffices to have

$$2^k - 1 > \lceil \frac{n}{k} \rceil$$

in order to be able to do the counting within fields of k bits each. For $n = 32$ we can choose $k = 4$ and obtain the following solution, which uses only 7 “C-operation” as opposed to the 8 operations of the “parity of word with a multiply”

from [1].

```
parity(x)
register unsigned x;
{
    x ^= x >> 1; // parity of 2 bit fields
    x ^= x >> 2; // parity of 4 bit fields
    x &= 0x11111111; // select low bits of 4 bit fields
    return((x % 0xf) & 0x1); // apply HAKMEM technique,
                             // return least sign. bit
}
```

This approach would work for word-lengths up to 56. Two more instructions would suffice for up to 2032 bits (with constants adjusted).

4 Discussion

We have obtained several bit counting and parity algorithms for restricted sets of operations. Asymptotically the (modified) parity function with complexity $O(\log n)$ for the OPAL-model is as fast as the known solutions based on broad-word steps [7]. The latter may include shift-operations by a constant number of bit positions. A lower time bound $\Omega(\log n / \log \log n)$ follows from a result in circuit complexity. The bound is tight for circuits, but it is open if this is also true for the models investigated here.

Appendix

HAKMEM item 169 [3] describes a reduction of bit counting for word-lengths of at most 62 bits to integer division. In its original form the program is presented in assembly language for the PDP-6/10 (36 bit architectures), which is not very well-known nowadays.¹

We render it here in C (where the variable names **a** and **b** correspond to the registers of the original program, comments are given separately):

```
#define TWOBITS 033333333333
#define THREEBITS 030707070707

bitcount(a)
unsigned a;
{
    unsigned b;
```

¹Most of the PDP-6/10 instructions are quite suggestive. Two possible exceptions are:

LDB B,[014300,,A]: Load 35 bits (octal 043) from A with an offset 1, counting bits from the right. Store them right adjusted in B (line 1 of the C program, this cannot be implemented directly in C and we use the approach mentioned in the comment of [3, item 169] on the LDB instruction).

SUBB A,B: Subtract and store in both A and B (lines 6 and 7 of the C program).

See [2] for more information on the PDP-6/10 instruction set.

```

    b = a>>1;          // line 1
    b &= TWOBITS;       // line 2
    a -= b;            // line 3
    b >>= 1;
    b &= TWOBITS;
    a -= b;            // line 6
    b = a;             // line 7
    b >>= 3;
    a += b;            // line 9
    a &= THREEBITS;
    return(a % 077); // line 11
}

```

Comments:

lines 1-6 : Consider an octal digit of variable **a** composed of three bits x , y , and z with weight $4x + 2y + z$. Then in line 3 the value $2x + y + z$ is computed and in line 6 the sum $x + y + z$ of the three original bits is computed.

line 7 : This simulates the second transfer of SUBB.

line 9 : Neighboring groups of octal digits are added in order to compute the sum of six bits. Notice that the three resulting bits are able to hold the maximum count for 6 bits without a carry to the next octal digit.

line 11 : Consider the contents of **a** as a number in base 64 representation. Then the digit at position i from the right (starting at 0) has weight

$$64^i = (63 + 1)^i = 63^i + 63^{i-1}i + \cdots + 63i + 1$$

and modulo 63 the sum of all base 64 digits is computed. This limits the admissible word-length to 62 bits.

References

- [1] S. E. Anderson. Bit twiddling hacks. <http://graphics.stanford.edu/~seander/bithacks.html>.
- [2] H. Baker. PDP-10 machine language. <http://pdp10.nocrew.org/docs/instruction-set>.
- [3] M. Beeler, R. W. Gosper, and R. Schroepfel. Hakmem. Technical Report AIM-239, Massachusetts Institute of Technology, Cambridge, MA, USA, 1972.
- [4] S. Brumme. the bit twiddler. <http://bits.stephan-brumme.com/>.
- [5] R. J. Fisher and H. G. Dietz. Compiling for SIMD within a register. In *11th Annual Workshop on Languages and Compilers for Parallel Computing (LCPC98)*, pages 290–304. Springer Verlag, Chapel Hill, 1998.
- [6] B. W. Kernighan and D. Ritchie. *The C Programming Language*. Prentice-Hall, 1978.

- [7] D. Knuth. *The Art of Computer Programming: Vol 4, Fascicles 0-4*. Prentice Hall, 2009.
- [8] H. S. Warren Jr. Functions realizable with word-parallel logical and two's-complement addition instructions. *Commun. ACM*, 20(6):439–441, 1977.
- [9] P. Wegner. A technique for counting ones in a binary computer. *Commun. ACM*, 3(5):322, 1960.